Evolution of the Female Labour Force Participation Rate in Canada, 1976-1994: A Cohort Analysis

W-99-4E

by
Paul Beaudry and Thomas Lemieux

January 1999

Canada

Applied Research Branch
Strategic Policy
Human Resources Development Canada

Direction générale de la recherche appliquée
Politique stratégique
Développement des ressources humaines Canada

Evolution of the Female Labour Force Participation Rate in Canada, 1976-1994: A Cohort Analysis

W-99-4E

by
Paul Beaudry and Thomas Lemieux
January 1999

■

The Working Paper Series includes analytical studies and research conducted under the auspices of the Applied Research Branch of Strategic Policy. Papers published in this series incorporate primary research with an empirical or original conceptual orientation, generally forming part of a broader or longer-term program of research in progress. Readers of the series are encouraged to contact the authors with comments and suggestions.

La série des documents de travail comprend des études analytiques et des travaux de recherche réalisés sous l'égide de la Direction générale de la recherche appliquée, Politique stratégique. Il s'agit notamment de recherches primaires, soit empiriques ou originales et parfois conceptuelles, généralement menées dans le cadre d'un programme de recherche plus vaste ou de plus longue durée. Les lecteurs de cette série sont encouragés à faire part de leurs observations et de leurs suggestions aux auteurs.

Printed/Imprimé 2000
ISBN: 0-662-28205-1
Cat. No./N° de cat. MP 32-28/99-4E

General enquiries regarding the documents published by the Applied Research Branch should be addressed to:

Publications Office
Applied Research Branch
Strategic Policy
Human Resources Development Canada
165 Hôtel de Ville Street, Phase II, 7th Floor
Hull, Quebec, Canada
K1A 0J2

Telephone: (819) 994-3304
Facsimile: (819) 953-9077
E-mail: research@spg.org
http://www.hrdc-drhc.gc.ca/arb/

Si vous avez des questions concernant les documents publiés par la Direction générale de la recherche appliquée, veuillez communiquer avec :

Service des publications
Direction générale de la recherche appliquée
Politique stratégique
Développement des ressources humaines Canada
165, rue Hôtel de Ville, Phase II, 7e étage
Hull (Québec) Canada
K1A 0J2

Téléphone : (819) 994-3304
Télécopieur : (819) 953-9077
Courrier électronique : research@spg.org
http://www.hrdc-drhc.gc.ca/dgra/

Abstract

This paper assesses the contribution of cohort effects, age composition effects, and macroeconomic factors on the evolution of the female labour force participation rate in Canada between 1976 and 1994. Using data from the Survey of Consumer Finances, the authors find that cohort effects are the main factor behind the recent stagnation in female participation rates. Though the poor macroeconomic performance of the Canadian economy during the 1990s has also contributed to this stagnation phenomenon, it cannot explain in itself why the behaviour of female participation rates in the 1990s was so different from that of previous decades. The authors reach similar conclusions when analyzing the evolution of the employment rate.

One related finding is that both the level and the slope of the age-participation profiles of women have changed over time. While older cohorts had profiles that were sloping up between the ages of 25 and 50, younger cohorts exhibit much flatter (and higher) profiles for the same age range. In other words, age-participation profiles of women increasingly resemble those of men which are flat at very high levels before declining after age 50.

Résumé

Dans ce texte, les auteurs évaluent le rôle des effets de cohortes, de la structure d'âge et des facteurs macro-économiques dans l'évolution du taux d'activité des femmes au Canada entre 1976 et 1994. Le résultat principal qui ressort de l'analyse des données de l'enquête des finances des consommateurs est que les effets de cohortes sont le facteur clef permettant d'expliquer le plafonnement récent de ces taux d'activité. Bien que la performance macro-économique peu enviable de l'économie canadienne dans les années quatre-vingt-dix ait elle aussi contribué à ce phénomène de plafonnement, il n'en reste pas moins que seuls les effets de cohortes parviennent à expliquer pourquoi les taux d'activité ont crû beaucoup plus rapidement dans les années soixante-dix et quatre-vingt que dans les années quatre-vingt-dix. Les auteurs tirent les mêmes conclusions lorsqu'ils analysent les taux d'emploi plutôt que les taux d'activité.

Il ressort aussi de l'étude qu'à la fois le niveau et la pente des profils de participation en fonction de l'âge (profils d'âge) ont changé à travers le temps. Alors que la pente des profils d'âge était positive entre l'âge de 25 et 50 ans pour les cohortes entrées depuis longtemps sur le marché du travail, ces profils sont beaucoup plus plats pour celle entrées plus récemment sur le marché. En d'autres termes, les profils d'âge des femmes ressemblent de plus en plus à ceux des hommes qui sont, eux aussi, plutôt plats jusqu'à l'âge de 50 ans.

Acknowledgements

We thank Richard Archambault, Allan Crawford, Pierre Fortin, David Green, Louis Grignon, Andrew Sharpe, and two anonymous referees for their comments. We also appreciate the help of David Green, Maud-Catherine Rivard and Alexandre Debs with the data.

Table of Contents

1. Introduction

One striking feature of the Canadian labour market during the 1970s and 1980s was the sustained growth in the labour force participation rate of women. Figure 1a shows that for women age 25 to 64, the participation rate increased from less than 50 percent in the mid 1970s to 70 percent in the late 1980s. An equally striking development is the recent stagnation in the female participation rate that has remained around 70 percent throughout the 1990s. While the participation rate of men of the same age group has declined during this latter period, Figure 1a shows that the 1990s represented a much less dramatic departure from previous trends—indicated by a dotted line in the figure—for men than for women.

The goal of this paper is to investigate the sources of the stagnation in the labour force participation rate of women age 25 to 64 during the 1990s. More precisely, we seek to determine whether the stagnation in the female labour force participation rate is a temporary phenomenon tied to the poor growth performance of the economy during the 1990s, or whether it represents a permanent change in the behaviour of women in the labour market. For example, the stagnation in the participation rate may signal that the process of integration of women into the work force is almost complete. A cohort analysis is performed using data from the Survey of Consumer Finances (1976 to 1994) to examine this issue.

The methodology involves isolating the effect of three separate factors on the participation rate of women. To this end, we follow cohorts of women over time, that is, we track the participation rate of representative groups of women who entered into the work force at a given point in time (e.g. women who were 25 years old in 1976). We then decompose a cohort's participation rate into, first, a *macroeconomic effect* that by definition is common across cohorts. Recession and structural phenomena such as the generosity of the employment insurance system are some of the factors that may cause a macroeconomic effect. The second factor is the *age* or *life cycle effect*, which shows how the cohort's participation rate changes as the cohort ages. The third factor is the *cohort specific effect*, which shows differences between cohorts for a given age and macroeconomic effect. For example, if the cohort that entered the labour force in 1976 has a participation rate that is 10 percent higher than that of the cohort that entered the labour force in

1966 at the same age and under similar macroeconomic conditions, the 1976 cohort is said to exhibit a 10 percent cohort effect relative to the 1966 cohort.

Our results indicate that cohort effects are likely the dominant factor in explaining the recent stagnation in female labour force participation rates. The same result was obtained when the labour force participation rate was replaced by the employment rate. Cohort effects help explain both the large increase in participation and employment rates during the 1970s and 1980s, as well as their stagnation in the 1990s. The 1989-1994 recession merely amplified the stagnation phenomenon; it also explains the observed decline in the participation rate for some demographic groups. These results show, however, that stagnation would have occurred, albeit later in the 1990s, even if more favourable macroeconomic conditions had prevailed.

This paper is divided into seven sections. Section 2 outlines the data used in this study. Section 3 presents the actual cohort analyses. In Section 4, the results of those analyses are used to break down the evolution of the labour force participation rate into macroeconomic, age and cohort effects. Section 5 examines the robustness of the results, studying both the influence of changes in the population's education level on changes in the participation and employment rates, and the effect of the generosity of the employment insurance program and of the real wage rate of women. And finally, Section 6 makes some projections for future participation and employment rates. Participation and employment rates will be analyzed simultaneously to ascertain that the evolution of the participation rate does not merely reflect changes in the way Canadian women "classify" themselves in the labour market.[1] The seventh section concludes the paper.

[1] Card and Riddell's results (1993) suggest that the increase in Canada's unemployment rate (relative to the U.S.) in the 1990s can be explained in part by how individuals answer questions on labour force participation on the Labour Force Survey (LFS).

2. Data and Descriptive Statistics

The data used were obtained from the Survey of Consumer Finances (SCF) for the years 1976, 1978, 1980, 1982, 1983, 1985, 1987, 1989, 1991, 1993 and 1994 (survey years).[2] These years were used because: 1) from 1976 to 1982, the survey was taken only every other year; 2) since then, the survey has been taken every year except 1984; and 3) 1994 was the last year available when we began this study. The years represented (about every other year) provide a fairly coherent sample across time. One important advantage of the SCF over other existing data sets is that the age of respondents is available in the public use samples provided by Statistics Canada.[3]

Individuals were grouped into two-year cohorts according to their date of permanent entry into the labour force. This was defined, somewhat arbitrarily, as the even-numbered year in which the woman in question was 25 or 26 years old (e.g. a woman born in 1941 is in the "entering" cohort for 1966). Thus for each even year, all women between 25 and 64 are divided into 20 cohorts (25-26, 27-28, ..., 63-64).[4] In total, 29 cohorts entered into the labour force between 1936 and 1992.

Note that Statistics Canada's public use files of the SCF from before 1982 provide data only for heads of households and spouses; we therefore confined our analysis to this sub-sample for the entire 1976-1994 period.[5] Labour force activity (employment, unemployment or non-participation) is determined based on individual responses to the usual LFS questions (for the month of April in the SCF). Figure 2 shows the evolution of the labour force participation rate, represented by a solid line, and the employment rate, represented by a dotted line, for the cohorts which entered the labour force in 1940, 1950, 1960, 1970, 1980 and 1990. These six cohorts are

[2] SCF data are also available for 1972 to 1975. We are not using them, however, due to a problem of non-comparability of the questions on participation and employment, which were revamped as part of major changes to the Labour Force Survey in 1976.

[3] The public use samples of the Labour Force survey have bigger sample sizes and are available for every year since 1976. Unfortunately, this data set in unsuitable for a detailed cohort analysis because respondents are pooled into 5 year age groups.

[4] For odd-numbered years (1983 to 1993), the women included are between 25 and 65 years old.

[5] Beaudry and Green (1996) point out that this restriction has little impact on individuals aged 25 and older, who are almost all heads of households or wives.

a relatively representative sample of the full 29 cohorts that are not shown to simplify the figure. The cohorts that entered the labour force first (e.g. the 1940 cohort) appear only in their last years in the labour market, while those that entered last (e.g. the 1990 cohort) appear only in their earlier years in the market. The figure shows that the evolution of the labour force participation and employment rates is similar for all cohorts. Both these rates tend to increase from the age of 25 to 45-50 years, then decrease rapidly until age 65. Participation and employment rates are obviously higher for those cohorts that entered the labour force most recently than for the others.

Other descriptive statistics are presented in Table 1, which illustrates age composition and education levels (percentage of women with a high school education or less) for each of the years studied. The table shows quite a young population during the period from 1976 to 1994. About 65 percent of women between the ages of 25 and 64 during these years were 44 or younger.

Table 1: Descriptive Statistics

Year	Age distribution in percentages				Percentage with high school or less	Number of observations
	25-34	35-44	45-54	55-64		
1976	40.9	24	19.8	15.3	73.6	12,269
1978	39.1	24.9	20.5	15.5	75	17,372
1980	39.9	26.1	19.1	14.9	74.7	18,212
1982	39.5	26.2	19.1	15.3	73.5	18,881
1983	39.2	27.5	18.9	14.4	72.7	19,775
1985	38.7	29.3	17.7	14.3	71	19,664
1987	37.9	29.9	18	14.3	70.4	17,949
1989	37.1	31	18.3	13.6	68.3	21,117
1991	34.9	32.3	19.3	13.6	59.7	26,033
1993	31.4	32.8	21.6	14.2	56.2	22,592
1994	32.8	32.7	21.2	13.3	54.3	22,420
Total	37	29.2	19.4	14.3	67.2	216,284

The impact of the baby boom/baby bust on the population's age composition is also clearly visible. This helps explain the increase in the proportion of women aged 35-44 since the beginning of the 1980s; the first wave of boomers born in 1946 reached the age of 35 in 1981. The same phenomenon occurred at the beginning of the 1990s as the first of the boomers reached

45. It is now the baby bust generation, those women born after 1965, who make up the 25-34 year-old segment.

The statistics presented in Table 1 also show a steady increase in level of education: the percentage of women with a high school education or less dropped from 73.6 percent in 1976 to 54.3 percent in 1994. This trend, however, is slightly exaggerated by the revamping of the questions on education in the LFS in 1990.[6]

[6] Before 1990, LFS questions on education did not specifically mention trade certificates. Since that time, those who have finished high school and also hold a trade certificate have been placed in the "post-secondary" category. This explains the break in the series between 1989 and 1991. Note also that changes to the LFS questionnaire make it difficult to use a more detailed classification than "high school or less" versus "post-secondary."

3. Cohort Analysis

3.1 Econometric Model

An econometric model is used to examine the separate roles played by the macroeconomic, cohort and age effects on labour force participation and employment rates. The dependent variable used in the regressions is the participation (or employment) rate p_{jt} for cohort j at time t expressed in "log-odds" form $ln(p_{jt}/(1-p_{jt}))$.[7] For example, $p_{74,84}$ represents the labour force participation rate for the cohort that entered the labour force in 1974 ($j = 74$) during the year 1984 ($t = 84$). This functional form is used to account for the special nature of variable p_{jt}, whose value is always between 0 and 1. It ensures that the predicted value will always be between 0 and 1, which would not be the case if a standard linear specification were used instead.[8]

In most of the estimated models, only one macroeconomic variable is used, the unemployment rate among men aged 25 to 44. Although certain long-term trends in this rate may be determined by structural factors, it is clear that its short-term fluctuations are mainly reflective of the evolution of the economic climate. Other variables such as the output gap may be used in addition to unemployment rate, but we prefer to concentrate on the latter, because of its simplicity; however, the results must be interpreted with caution.[9] The scope of the macroeconomic effect will, however, be broadened in Section 5 by adding other variables.

The model is completed by adding a flexible specification for age, or life cycle, effects (fourth degree polynomial) and for cohort effects (third degree polynomial). This yields the equation:

$$(1) \qquad ln(p_{jt}/(1-p_{jt})) = \alpha + \delta ur_t + \beta_1 j + \beta_2 j^2 + \beta_3 j^3 + \gamma_1 a_{jt} + \gamma_2 a_{jt}^2 + \gamma_3 a_{jt}^3 + \gamma_4 a_{jt}^4$$

where ur_t represents the unemployment rate among men between 25 and 44, j represents the cohort, and a_{jt} represents the age of women from cohort j in year t.

---

[7] The "log-odds" specification is in fact only an application of the "logit" model to grouped data (grouped by cohort).

[8] Consider $q = ln(p/(1-p))$. This results in $p = \exp(q)/(1 + \exp(q))$ where $0 \leq p \leq 1$, since $\exp(q) \geq 0$. Thus, no matter what the value of q predicted by a regression such as equation (1), the predicted value of p will therefore always be between 0 and 1.

[9] In theory, the output gap should be more representative of the economic climate than unemployment among men aged 25 to 44. In fact, the validity of the output gap depends on the accuracy of a number of difficult-to-verify hypotheses on the structure of the macroeconomic model used to obtain this measurement.

One characteristic of equation (1) is that the age profile for each cohort, i.e. the evolution of the labour force participation rate over the life cycle, is similar for each cohort; they differ only in terms of the intercept. In other words, the model allows a vertical displacement of the life cycle profile from one cohort to another while forcing the shape of the profile, and thus the slope, to be identical for each cohort. A more general model is produced by introducing age-cohort interaction terms to allow the age effect to vary from one cohort to another. This was done with the following model, which incorporates an age-cohort ($a_{jt}j$) and an age-cohort squared interaction term.[10]

$$(2) \qquad ln(p_{jt}/(1-p_{jt})) = \alpha + \delta ur_t + \beta_1 j + \beta_2 j^2 + \gamma_1 a_{jt} + \gamma_2 a_{jt}^2 + \gamma_3 a_{jt}^3 + \gamma_4 a_{jt}^4 + \theta_1 a_{jt} j + \theta_2 a_{jt} j^2.$$

If second or higher order polynomial terms are omitted, equation (2) shows that the effect of age on $ln(p_{jt}/(1-p_{jt}))$ is equal to $\gamma_1 + \theta_1 j$. If θ_1 is positive, the age effect will be greater for those cohorts which most recently entered the labour force (highest j) than for the others, and vice versa. Coefficient θ_1 thus allows the life cycle profile to vary from one cohort to another.

Graphs 1a and 1b provide illustrative examples of the advantages of equation (2) over equation (1), which does not include the age-cohort interaction term. Without such interaction terms, the intercept is the only difference between different cohorts' age profiles (Graph 1a). The same increase in participation at career outset and the same decrease in participation at career end is shown for every cohort. The age profile is clearly more flexible in Graph 1b where interaction terms are introduced. In this graph, the "newer" cohort has both a higher ordinate value at the origin and a shallower slope. This results in a higher and more stable age profile at career outset than in the previous cohorts (the "older" cohorts on the graph). This profile is also more similar to that for men, whose participation rates are fairly high and stable until the age of about fifty. The situation shown in Graph 1b is therefore more consistent with the idea of a convergence between men's and women's labour force participation rates, or increasing participation of women in the labour force, than that shown in Graph 1a.

In Graph 1b, the cohort effect is concentrated at career outset, participation rates before the age of 40 for the newer cohort being much higher than those for the older cohort, while the rates are

[10] This particular functional form was arrived at using specification tests. As a rule, we add terms of a higher and higher power until the terms added are no longer significant. The same procedure was used for equation (1).

reasonably comparable after the age of 50. The impact of the entry of the newer cohort on the aggregate labour participation rate would thus be felt most strongly during the first 10 or 20 years after the cohort's arrival, while in Graph 1a, its influence is shown as continuing throughout the life cycle. In other words, the entry of newer cohorts in Graph 1b should result in a rapid increase in the aggregate participation rate, followed by a period of stagnation. Graph 1a, on the other hand, implies a constant increase in the aggregate participation rate.

Graph 2 illustrates the impact of the arrival of newer cohorts in the two cases discussed above, with those cohorts entering the labour force after 1970 considered "newer" cohorts and those entering before 1970 considered "older" cohorts.[11] The graph clearly shows that only the presence of an age-cohort interaction effect can explain the stagnation phenomenon.

Though the case presented in Graphs 1a, 1b and 2 are only illustrative examples, they provide the intuition on the importance of the age-cohort interaction in explaining the stagnation of participation rates in the 1990s. The specification we actually estimate (equation 2) is a flexible version of the dichotomous case considered in Graphs 1a, 1b and 2. It allows for smooth changes in age profiles across cohorts.[12]

An often-mentioned problem with cohort analyses is the impossibility of separately identifying cohort effects, year effects (macroeconomic effects), and age effects because of the linear dependence between them. In fact, since $a_{jt}=25+t-j$, the three variables (a_{jt}, j and t) are perfectly collinear. This study proceeds on the implicit assumption that variable ur_t captures any systematic macroeconomic effect and that there is no other temporal trend in this effect. If their was a temporal trend in the female participation rate, due for instance to the feminist movement, we would be attributing this trend to cohort effects. Note, however, that as long as the temporal trend is linear, it cannot account in itself for the stagnation in participation rates (a break in trend) in the 1990s. So although there may be some issues regarding the interpretation of our results,

[11] We also assume that age composition is uniform (one fortieth of the population at each age).

[12] The specification proposed in equation (2) fits the data much better than a specification more directly in the spirit of the illustrative example (an "older" and a "newer" cohort with completely separate age profiles). The problem with the illustrative example is that it implies a sharp discontinuity across cohorts which is inconsistent with the observed data. Equation (2) also implies some restrictions since only the linear term of the age profile is allowed to vary across cohorts. We chose that latter specification because allowing for more general interaction terms did not affect our main conclusions.

we do not think that our analysis of the sources of the stagnation in female participation rates will be affected by these issues.

These considerations aside, econometric models (such as 1 and 2) can never explain all the variations in the data (R squared is less than 1). As a rule, a *residual macroeconomic effect* is obtained, representing the macroeconomic variation in the data that cannot be explained by other variables in the model. If during a period, say the 1990s, we were to find a large residual we would interpret this as indicating that participation in this period has experienced a macroeconomic effect not captured by its standard co-movement with the unemployment rate.

3.2 Results

Equations (1) and (2) were estimated using weighted least squares, with cohort size j at time t used as the weights. The results are shown in Table 2 for employment rates (columns 1 to 3) and participation rates (columns 4 to 6). For model 1, note that all the coefficients are significant except for the rate of unemployment among men 25 to 44 (columns 1 and 4). That effect becomes significant, however, when age-cohort interaction terms from model 2 are introduced (columns 2 and 5). Also note that interaction term coefficients are highly significant, and that R-squared for model 2 is higher than for model 1.

We also present the results of the regressions when the sample is limited to the 1976-1989 period (columns 3 and 6). The purpose of this exercise is to assess whether the leveling off of participation and employment rates in the 1990s was predictable from the behaviour of these rates prior to 1990. The results indicate that the estimated parameters for 1976-1989 are relatively similar to those for the period as a whole. In fact, the p-value of the Chow-test statistic reported at the bottom of column 2 indicates that we cannot reject the null hypothesis that the estimated employment rate models are the same for 1976-94 (column 2) and for 1976-89 (column 3).[13] The estimated models for 1976-94 and 1976-89 are only marginally different (p-value of 0.036 in column 5) in the case of the participation rate. We shall return to the question of the stagnation of the participation and employment rates.

[13] This version of the Chow test compares the estimated models for a prediction sample (1976 to 1989) and for the full sample (1976 to 1994). It can either be interpreted as a model stability test (Are the estimated models the same for the two sample periods?) or as a prediction test (Does the model estimated for the prediction sample predict accurately for other years?).

Table 2: Detailed Results of Regression
(Standard deviation in brackets)

	Employment Rate			Participation Rate		
	1976-1994	1976-1994	1976-1989	1976-1994	1976-1994	1976-1989
	1	2	3	4	5	6
Constant	0.792	-3.102	-2.912	0.997	-3.398	-2.746
	(0.143)	(0.291)	(0.344)	(0.158)	(0.245)	(0.177)
Unemployment rate	-0.019	-0.029	-0.031	-0.004	-0.016	-0.014
	(0.011)	(0.004)	(0.003)	(0.011)	(0.003)	(0.002)
Cohort effect:[a]						
co	-1.830	3.804	3.506	-2.227	4.162	3.267
	(0.154)	(0.383)	(0.465)	(0.151)	(0.319)	(0.215)
co^2	1.609	-0.792	-0.681	1.892	-0.851	-0.561
	(0.131)	(0.118)	(0.150)	(0.102)	(0.100)	(0.062)
co^3	-0.289	---	---	-0.336	---	---
	(0.025)			(0.019)		
Age effect:						
Age/10	0.389	2.706	2.780	0.405	3.004	2.749
	(0.053)	(0.175)	(0.216)	(0.051)	(0.124)	(0.126)
$Age/10^2$	-0.445	-1.742	-0.676	-0.471	-0.760	-0.701
	(0.024)	(0.101)	(0.027)	(0.021)	(0.032)	(0.019)
$Age/10^3$	-0.130	-0.151	-0.161	-0.139	-1.62	-0.167
	(0.007)	(0.007)	(0.009)	(0.006)	(0.006)	(0.009)
$Age/10^4$	0.031	0.032	0.024	0.032	0.034	0.027
	(0.006)	(0.006)	(0.004)	(0.005)	(0.005)	(0.005)
Interactions:						
Age/10 *co	---	-1.742	-1.902	---	-1.963	-1.886
		(0.101)	(0.140)		(0.067)	(0.091)
Age/10 $*co^2$	---	0.226	0.298	---	0.258	0.305
		(0.021)	(0.021)		(0.016)	(0.021)
P-value of pred. test [b]	0.001	0.096	---	0	0.036	---
R squared	0.941	0.954	0.951	0.953	0.964	0.962
Number of observations	224	224	164	224	224	164

[a] co = (year of entry into the labout force – 1932) /20
[b] p-value of a Chow test that the 1990-94 data is correction precited by the model estimated from the 1976-1989 data

To facilitate presentation of the results, it is simpler to use a graphical approach than to examine the numbers presented in Table 2 in detail. For each rate (participation and employment) and each model (1 and 2), we present the following four graphs: Graph (a) shows the cohort effect at age 44—i.e. the variations in the participation and employment rates attributable to the cohort effect at a precise point in the life cycle. Graph (b) shows the age effect throughout the life cycle for a typical cohort (the one which entered the labour force in 1964). Graph (c) presents a similar result for the six representative cohorts of Figure 2 (the ones which entered the labour force in 1940, 1950, 1960, 1970, 1980 and 1990) to illustrate cohort differences over the entire age profile. It should be noted, however, that Graph (b) shows a predicted age profile for the entire life cycle while Graph (c) shows the profile only for the ages at which the cohort in question is observed in the data (1976 to 1994). Finally, Graph (d) indicates the degree to which the estimated model for 1976-1989 can be used to predict participation and employment rates for the entire 1976-1994 period.

Rather than discussing each of the graphs in detail, we will confine ourselves to noting a few highlights:

- All the estimated models indicate a levelling off in cohort effects (Figures 3a to 6a); the participation and employment rates for the 1992 cohort are comparable to the ones for the cohorts which entered the labour force in the 1980s (or at least, the cohorts will all be comparable once they have reached age 44).

- The participation and employment rates peak around age 50 (Figures 3b to 6b).

- The younger cohorts have flatter age profiles (shallower slopes) early in their careers. This pattern is particularly pronounced for model 2, which includes age-cohort interaction (Figures 4c and 6c).[14]

- The model without interactions (model 1) provides no explanation whatsoever of the stagnant participation rate (Figure 3d) and falling employment rate (Figure 5d) observed in the 1990s.

[14] Note that the slope of the age profile in the participation rate (or employment rate) and age space may change even without interaction terms due to the log-odds specification.

On the other hand, these phenomena can be predicted more accurately from the model with interactions (model 2) estimated for 1976-1989 (figures 4d and 6d).

This last finding is the most interesting. The figures suggest that model 2 does a much better job predicting post-1990 patterns from pre-1990 data than model 1. The prediction tests reported at the bottom of Table 2 is a formal way of evaluating this hypothesis. The p-value for the specifications corresponding to model 1 (columns 1 and 4) are well below the critical value of 0.05, which suggests that the models fit to the 1976-89 period *do not* predict accurately the post-1990 data. By contrast, the specifications corresponding to model 2 predict more accurately the post-1990 data, especially for the employment rate model (p-value of 0.096). Though the p-value is higher in column 5 (model 2) than in column 4 (model 1), we nevertheless reject the null hypothesis of an accurate prediction for the participation rate.

The reason why the prediction test fails for the participation rate can be seen in Figure 4d. The figure shows that the model predicts some growth in the participation rate in the early 1990s while the actual data show a clear stagnation.[15] Relative to the pre-1990 period for which the fit of the model is almost perfect (predicted and actual values look almost identical), the discrepancy between the predicted and actual values in the 1990s seems to be sizable.

All in all, however, model 2 provides a much more accurate prediction of what happened in the 1990s than model 1. Looking at the average for years 1991, 1992, and 1994 (three years of observation for the 1990s), model 1 overpredicts the employment and participation rates by 3.0 and 3.2 percentage points, respectively. By contrast, model 2 underpredicts the employment rate by 0.4 percentage points and overpredicts the participation rate by 0.5 percentage points.[16]

[15] The actual data shows a small increase in the participation rate from 1989 to 1991, followed by a small decrease from 1991 to 1993. However, these variations are not statistically significant. It is thus accurate to describe the period from 1989 on as one of stagnation in the participation rate. The stagnation in the participation rate can be seen more clearly in Figure 1a which is based on the larger and, therefore, more accurate samples of the Labour Force Survey.

[16] An alternative set of "prediction tests" consists of verifying whether these average prediction errors are significantly different from zero. These tests yield the same inference as the Chow tests reported in Table 2, namely that only model 2 for the employment rate predicts accurately the post-1990 data.

Our results seem to indicate that in addition to unfavourable macroeconomic conditions, the levelling off of cohort effects also contributed to the trend observed in the 1990s. This hypothesis will be examined in detail in the following section. Note also that the stable cohort effects for cohorts which entered after 1970 yields some support to the assumption that there is not temporal trend in participation rates (Section 3.1). If there was a positive temporal trend in participation rates, the cohort effects for cohorts which entered after 1970 would have to follow a negative trend (youngest cohorts less active than those which entered in 1970), which is a surprising result.[17]

[17] To see this point, consider the simple linear participation model $p=\alpha+\beta j+\gamma a$, where the time trend has been set arbitrarily to zero. Since $t=j+a-25$, if there is a true time trend δ, an equivalent specification can be rewritten as $p=\alpha+\beta j+\gamma a+\delta t-\delta t = \alpha+\beta j+\gamma a+\delta t-\delta(j+a-25) =(\alpha+25\delta)+(\beta-\delta)j+(\gamma-\delta)a+\delta t$. In other words, the cohort effect becomes $(\beta-\delta)$ instead of β. If β is equal to zero (stable cohort effects as for the post-1970 cohorts), this means that, controlling for age and temporal effects, participation rates are lower for younger (high j) than older (low j) cohorts, which is unappealing on *a priori* grounds.

4. Decomposition

We shall now perform a more formal analysis of the role of different factors in the recent evolution of aggregate participation and employment rates for all women aged 25-64 by decomposing this evolution into four components: the macroeconomic effect related to the unemployment rate among men aged 25-44 (the economic cycle), the residual macroeconomic effect, the age effect and cohort effects. In terms of equation (2), it is relatively easy to identify the first two factors, which correspond to the term δur_t and to the residuals of this equation.[18]

More precisely, we first calculate the participation (or employment) rate for each year, taking the weighted average of p_{jt} values for each t. The observed rate (p_{jt}) is then replaced by the predicted rate $\hat{p}_{jt}$ from the estimated model.[19] The average of the $\hat{p}_{jt}$ values for each t therefore represents the aggregate rate predicted by the model. The difference between the observed aggregate rate and the predicted rate represents the residual macroeconomic effect.

We then recalculate the prediction by replacing the observed unemployment rate by the average of the unemployment rates over the entire sample (8.2 %). The difference between this new prediction and the preceding prediction represents the macroeconomic effect related to the male unemployment rate, which we also call the *cyclical effect*.

The cyclical and residual effects obtained in this manner are presented in Figure 7a for the participation rate. Note that the figures showing the decomposition results for the employment rate are qualitatively similar to those for the participation rate and are not reported here.[20] We nevertheless discuss the main results both for the participation and employment rates in the text.

During the 1990s, the cyclical effect is in the order of -1% for the participation rate and -2% for the employment rate. In other words, the female participation rate would have been 1% higher in the first part of the 1990s if the male unemployment rate had held steady at 8.2%.

Age and cohort effects are somewhat more complicated to understand because of the interaction terms in equation (2). It should be noted, first of all, that the age effect comes into play only to

[18] The residual macro-economic effect is the average of the residuals for each year.

[19] The predicted rate (p) is derived from the predicted log-odd (l) by the following equation: $p=\exp(l)/(1+\exp(l))$.

[20] All these results are presented in a longer version of the paper (Beaudry and Lemieux, 1998).

the extent that the population's age composition changes over time. For example, the arrival of the baby boomers in the labour force in the early 1970s considerably rejuvenated the 25-64 year-old population as a whole. As these young women had below-average participation rates, it should have been expected that this change in composition would have had a negative effect on the aggregate participation rate and vice versa.

It might therefore be supposed that to identify the age effect, it is enough to recalculate the predicted rate using a uniform age composition (5% of the population aged 25-64 in each 2-year age group) instead of the observed age composition. The problem with this procedure is that it depends on the cohorts present in the labour force in each year, since the age profile is dependent on the cohorts through the interaction terms. This procedure therefore serves to isolate the age effect plus the crossed age-cohort effect.

The same problem arises when we want to isolate the role of cohorts. For example, we can try to recalculate the predicted rates by replacing the cohort effect expressed as $\beta_1 j + \beta_2 j^2 + \theta_1 a_{jt} j + \theta_2 a_{jt} j^2$ by the cohort effect obtained if the cohort is set at an arbitrary level such as j=70 ($\beta_1 70 + \beta_2 70^2 + \theta_1 a_{jt} 70 + \theta_2 a_{jt} 70^2$). This gives us the cohort effect plus the crossed age-cohort effect, in the same way as in the case of age. Once we have all this information, however, it is possible to separate the overall effect of age and cohorts into three components: a "pure" age effect that indicates how participation rates would evolve over time if all cohorts exhibited the same participation profile; a "pure" cohort effect that indicates how the participation rate for a population with a uniform age distribution would evolve over time as older cohorts get replaced by younger ones; a joint age-cohort that residual interactions between age and cohort effects.

These different effects are illustrated in Figures 7b, 7c and 7d. Let us take the example of Figure 7b: in this case, we use the cohort which entered the labour force in 1970 as a reference cohort for the decomposition. The cohort effect thus indicates the difference between the observed rates and the rates which would have prevailed had all the cohorts followed the same age profile as cohort 17, other factors being kept constant. This cohort effect is therefore the "pure" effect mentioned earlier. The graph also shows the "pure" age effect (for a given cohort and other factors) as well as the combined age and age-cohort effect (the age effect for the observed cohorts in each year).

While it can be rather difficult to grasp all the details of these decompositions, the results speak for themselves: it is really the cohort effect that dominates the evolution of the participation and employment rates throughout the 1976-1994 period. The results are very similar regardless of which cohort is used as a reference for the decompositions (1970, 1980 or 1990). We find that cohort effects account for an increase of about 20 percentage points in participation and employment rates between 1976 and 1994. At the same time, the figures clearly indicate that this phenomenon seems to be coming to an end. By comparison, age effects play a relatively small role in recent changes.

To sum up, our results indicate that the stagnation of female participation and employment rates is primarily a structural phenomenon related to the stabilization of the cohort effects which were responsible for the remarkable increase in these rates in the 1970s and 1980s. The unfavourable macroeconomic situation amplified this phenomenon but was not the root cause. The relative performance of the participation and employment rates during the 1981-1983 and 1989-1994 recessions clearly illustrates this phenomenon; in 1981-1983, the downward pressure on the rates from the macroeconomic effect was offset by the cohort effects, pushing the rates up by one percentage point per year, whereas in 1989-1994, due to the stabilization of cohort effects, macroeconomic effects comparable to those of 1981-1983 resulted in lower participation and employment rates.

To clarify the role of cohort effects, we illustrate their magnitude at age 24, 34, 44, 54 and 64 in Figure 8. Let us take for example Figure 8c, which shows the cohort effect at age 44 by year of entry into the labour force. The vertical line indicates the cohort which was aged 44 in 1994. The curve to the left of the line describes the evolution of cohort effects during the 1976-1994 period. The curve to the right shows the predicted evolution for the coming years (see Section 6 for more details).

The results indicate a general slowing trend for most of the ages under consideration, attributable to cohort effects. This is particularly true for the younger groups (ages 24 and 34), which explains why the levelling off and declining trend is more pronounced for the 25-44 age bracket than for the 45-64 bracket (Beaudry and Lemieux, 1998).

5. Robustness Analysis

5.1 Education

We re-estimated the models separately for women who have pursued post-secondary studies and those who have only a high school diploma or less. The highlights of the results are:[21]

- The decline in the participation and employment rates in the 1990s is more pronounced among poorly educated women than for the female population as a whole. The growth in employment and participation between 1976 and 1989 is also weaker within each education group than for the population as a whole. A significant portion of the rise in the rates for the population as a whole therefore seems to be attributable to the increase in average education levels.

- Similarly, the cohort effects exhibit a decline for the most recent cohorts in most cases. This result suggests that the average quality of cohorts is declining, since a high education level is a less selective characteristic than it was in the past.

- The younger cohorts have very high and very flat age profiles for women who pursued post-secondary studies. These profiles are very similar to the ones for men with the same education levels.

5.2 Employment Insurance

We have re-estimated the regressions presented in Table 2 when the employment insurance subsidy rate is also used as a macroeconomic variable.[22] The results are not very conclusive, since the effect on the employment rate is negative when we also control for the unemployment rate among men aged 25-44.[23] The decline in the subsidy rate during the 1980s should therefore have increased the employment rate instead of lowering it. The effect on the participation rate is

[21] See Beaudry and Lemieux (1998) for detailed results.

[22] The subsidy rate is the product of the replacement rate and the maximum number of weeks of eligibility for a worker who worked the minimum number of weeks required for eligibility, divided by the minimum number of weeks required for eligibility (qualification period). Since the weeks of eligibility and the length of the qualification period depend on the local unemployment rate, we use a national weighted average of the subsidy rates in each UI economic regions. The detailed regression results are reported in Beaudry and Lemieux (1998).

[23] EI may have a different effect on the participation rate of different groups of women. For example, married women may react differently to changes in the EI subsidy rate than single women. The estimated effect in Table 3 may be thought as an average affect of the subsidy rate for the different subgroups of the population. Since the overall effects were not very encouraging, we have not explored how the results differed by subgroups.

not significant. This being said, including the subsidy rate as a macroeconomic variable has little impact on the model's other coefficients. Our conclusions about the role of cohort effects versus macroeconomic effects during the 1990s therefore remain unchanged.

5.3 Alternative Measures of the Cycle

In another attempt to examine the robustness of our results, we have re-estimated our main models with an alternative cyclical variable, the after-tax real wage of women, instead of the unemployment rate of men age 25 to 44.[24] This variable did not perform as well as a measure of the cycle than the unemployment rate of men age 25 to 44. For example, the estimated effect of the after-tax real wage of women was typically not statistically significant when the unemployment rate of men was also included in the estimated models. The other parameters of the models were very similar to those reported in Table 2.

[24] This wage variable used is the logarithm of after-tax annual earnings of full-year full-time female workers.

6. Predictions

We will now attempt to predict the future evolution of the participation and employment rates under two different macroeconomic scenarios: a 8.2% unemployment rate for men aged 25-44 (the average over the 1976-1989 period) and a 6.6% unemployment rate for the same group (the 1989 level). To do so, we must make some assumptions about the cohorts which will enter the labour force after 1994. We first assume that the age profile of these cohorts will be similar to those of the last cohort observed (the one which entered the labour force in 1992). This is a natural assumption to make since our empirical results indicate that the cohort effects have been relatively stable for the cohorts which entered the labour force in the 1980s and early 1990s (for example, see Figure 4a).[25] Since the age distribution is relatively uniform for cohorts that will enter the labour force throughout the 1990s, we also assume that these new cohorts have the same size as the cohort which entered the labour force in 1992.[26]

The results of the simulations are presented in Figures 9 (participation rate) and 10 (employment rate). The conclusions are the same in both cases: large increases in the participation and employment rates are clearly a thing of the past; in the future, these rates can be expected to hold relatively stable. However, there is still room for a 2-3 percentage point increase in the rates if the macroeconomic situation continues to improve. It is illusory, though, to think that the rates could rise 5-10 percentage points during the next period of expansion as they did over the 1983-1989 period. The cohort effects which prevailed at the time are simply no longer present today.

[25] Since an entering cohort represents only a small fraction of the population, our results would not change very much if cohort effects were to increase again for the newer cohorts.

[26] Tabulations from the SCF indicate that the 1992 cohort (women who turned 25 and 26 in 1992) represents 0.85 percent of the female adult population (15 and more). The four newer cohorts used in the simulations (age 17-18, 19-20, 21-22, and 23-34 in 1992) represent on average 0.83 percent of the female adult population, which is very similar to the 1992 cohort.

7. Conclusion

This study's main finding is that the leveling off of female participation and employment rates is primarily a structural phenomenon related to the stabilization of the cohort effects which accounted for the remarkable increase in these rates in the 1970s and 1980s. The unfavourable macroeconomic situation has amplified this phenomenon but is not the root cause. The relative performance of the participation and employment rates during the 1981-1983 and 1989-1994 recessions clearly illustrates this phenomenon; in 1981-1983, the downward pressure on the rates from the macroeconomic effect was offset by the cohort effects, pushing the rates up by one percentage point per year, whereas in 1989-1994, due to the stabilization of cohort effects, macroeconomic effects comparable to those of 1981-1983 resulted in lower participation and employment rates.

This result is strongly dependent on the amount of flexibility used to capture cohort effects. It is essential that the age profile as a whole, and particularly its slope, be allowed to vary from one cohort to another. This makes it possible to accurately trace both the rise and the flattening of the employment and participation profiles by age. These phenomena are consistent with a convergence in the behaviour of men and women in the labour market: men exhibit very high and very flat (at least until age 55) employment and participation profiles over their life cycle. The profiles of recent female cohorts are therefore closer to those of men than to those of older female cohorts.

Finally, the recent evolution of participation rates in the United States seems to corroborate our findings: Figure 11 shows that, despite more favourable macroeconomic conditions in the United States than in Canada, the U.S. female participation rate grew much slower in the 1990s than in the 1970s and 1980s. By contrast, the participation rate of U.S. men in the 1990s followed closely the pre-1990 trends. The fact that there was a sharp break from previous trends both in Canada and in the United States suggests that the poor performance of the Canadian economy during the 1990s was not the main factor behind the stagnation in female participation rates during this period. This reinforces our conclusion on the importance of cohorts effects, rather than cyclical factors, as the main explanation for this phenomenon.

Graphs

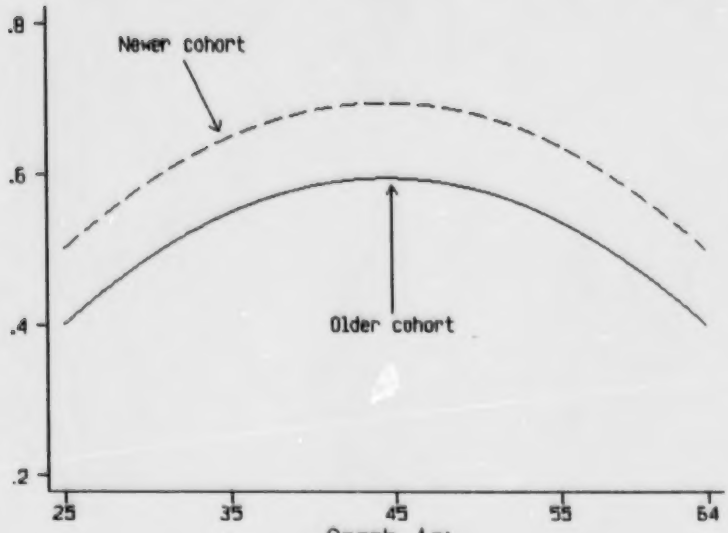

Graph 1a:
Age profiles without age-cohort interactions

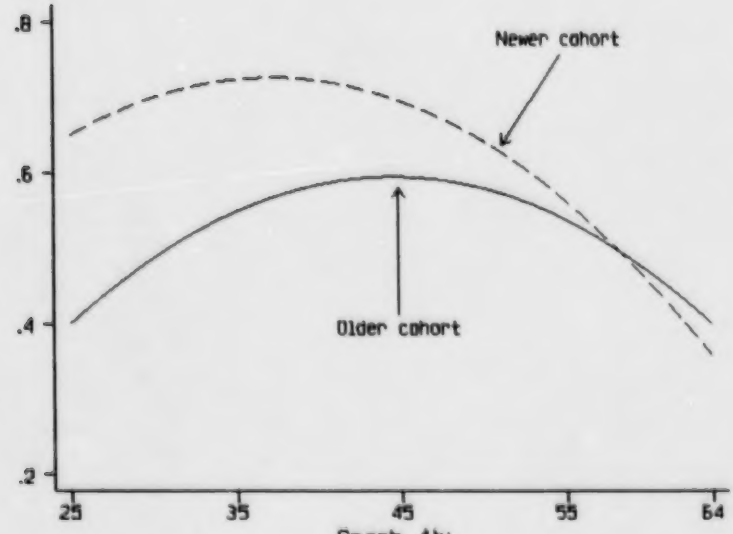

Graph 1b:
Age profiles with age-cohort interactions

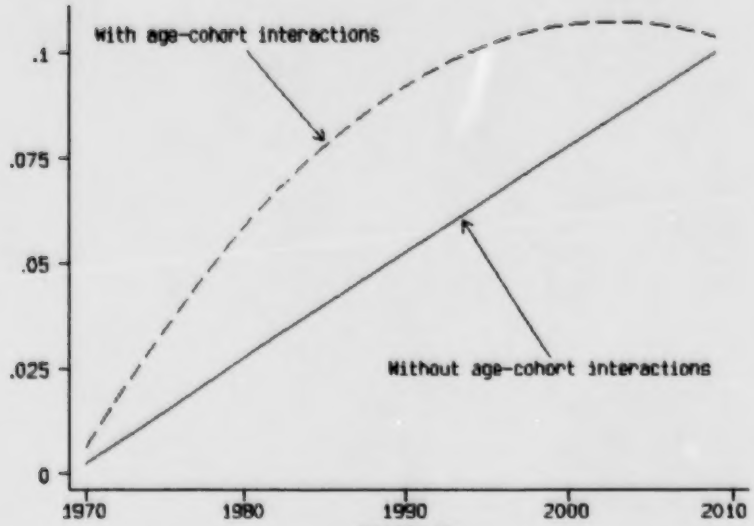

Graph 2:
Effect of the entry of new cohorts on the aggregate rate

Figures

Please note that the dotted line in all the figures represents the values of the participation rate (from 1990 to 1996) predicted by a linear trend model fit to the 1976 to 1989 period.

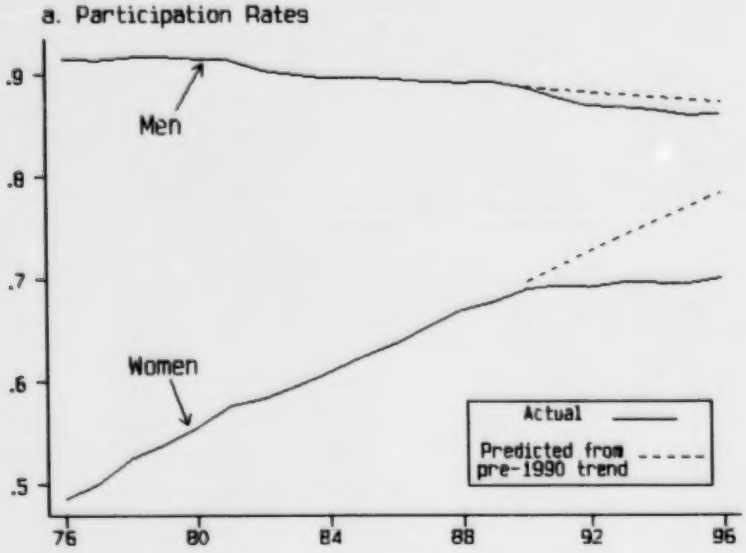

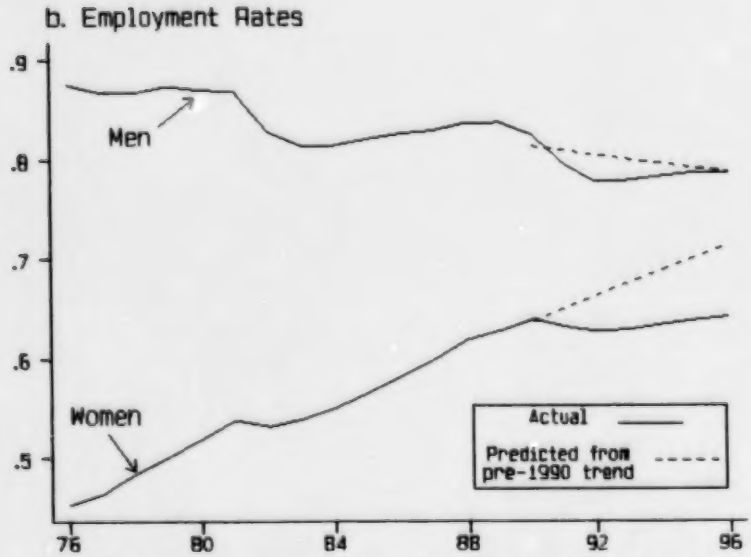

Figure 1: Participation and Employment Rates: Men and Women Age 25 to 64

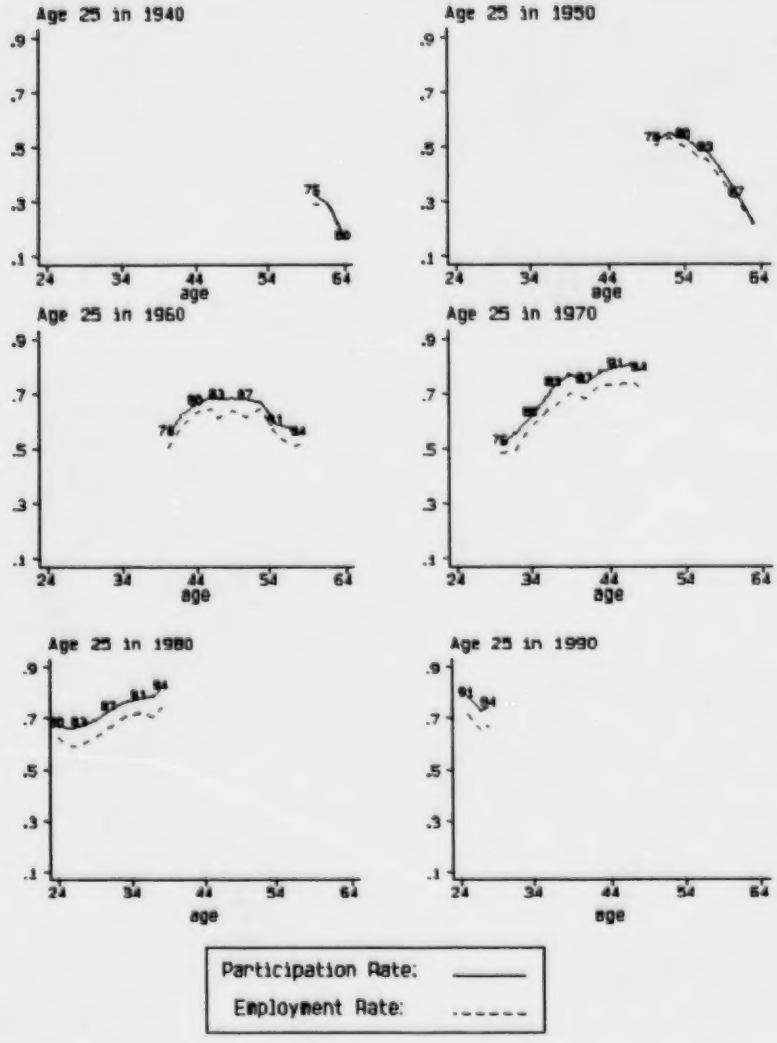

Figure 2: Participation Rate and Employment Rate by Cohort

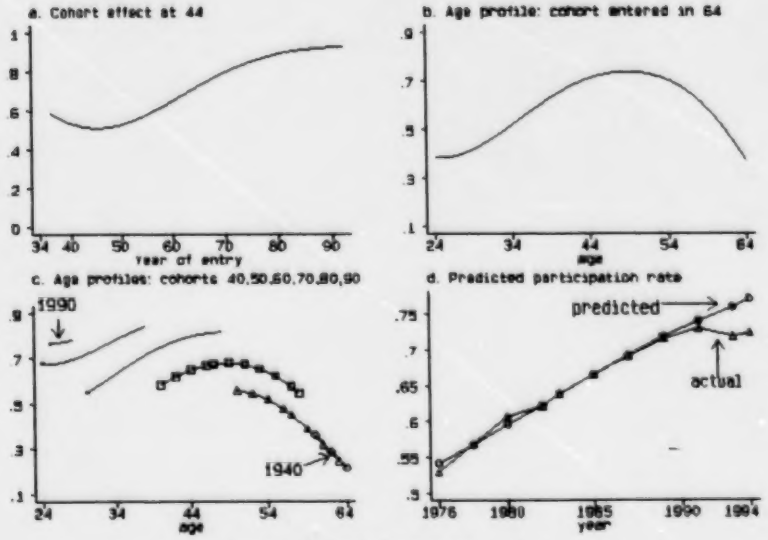

Figure 3: Participation rate
Model without age-cohort interactions

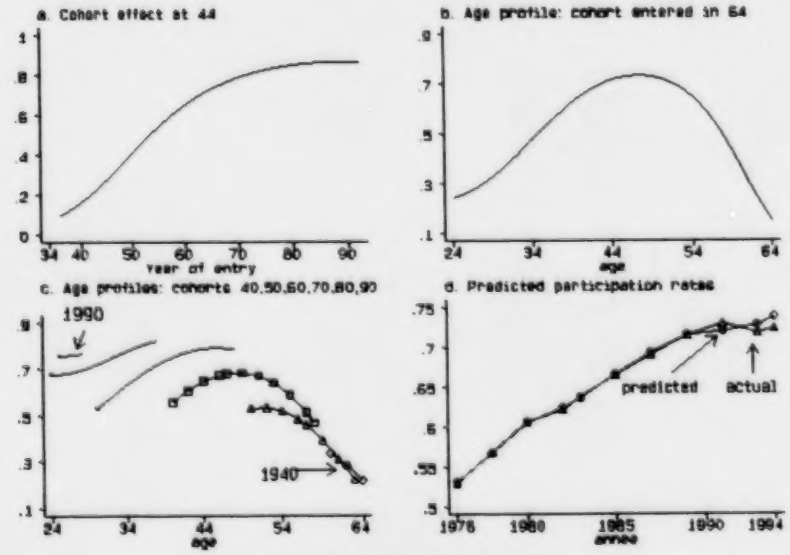

Figure 4: Participation rate
Model with age-cohort interactions

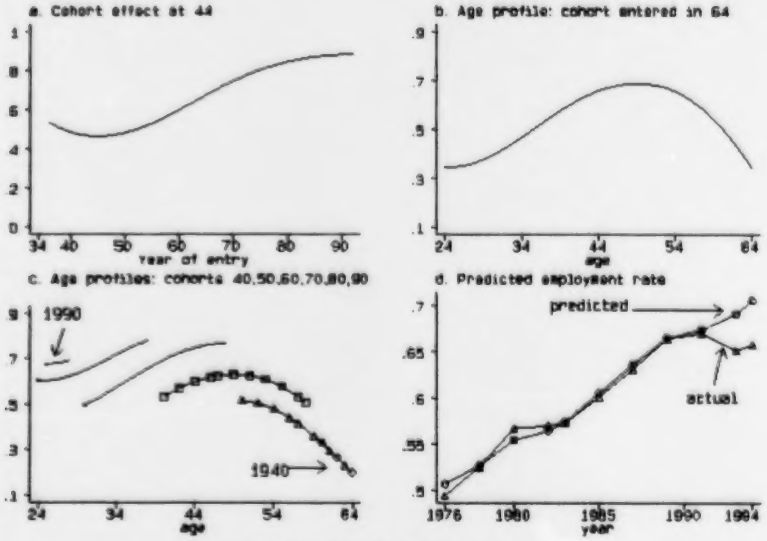

Figure 5: Employment rate
Model without age-cohort interactions

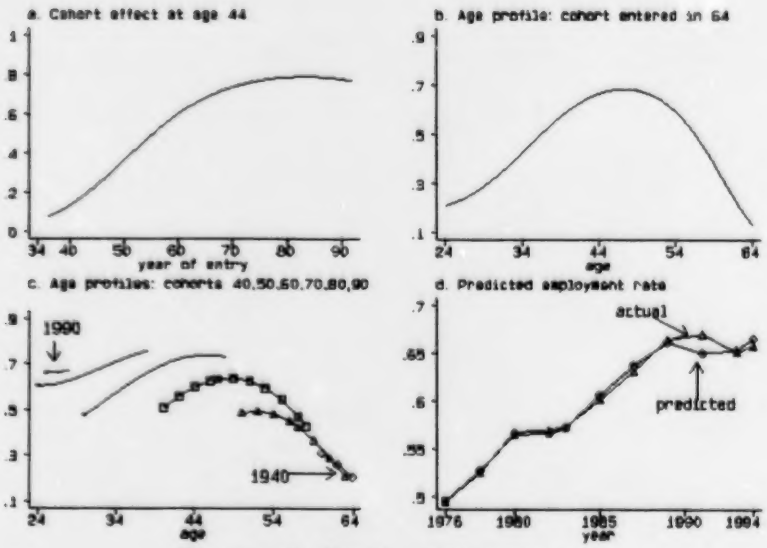

Figure 6: Employment rate
Model with age-cohort interactions

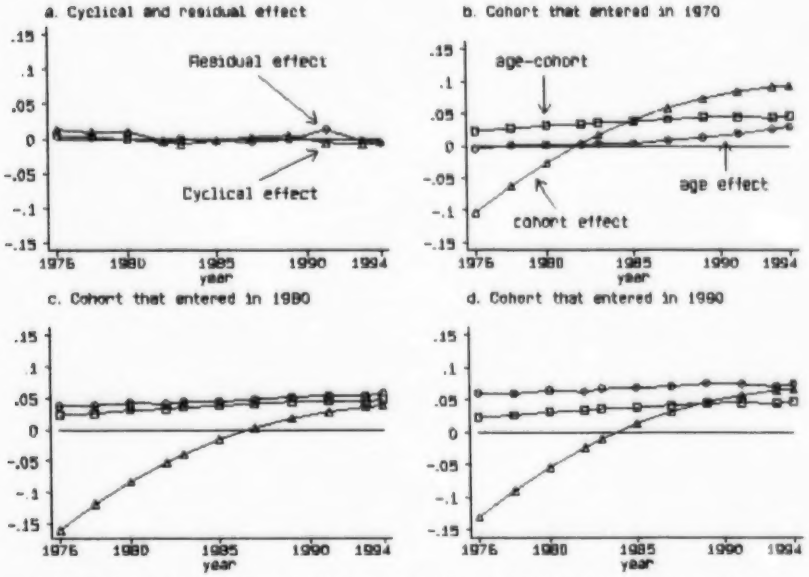

Figure 7: Decomposition of the participation rate

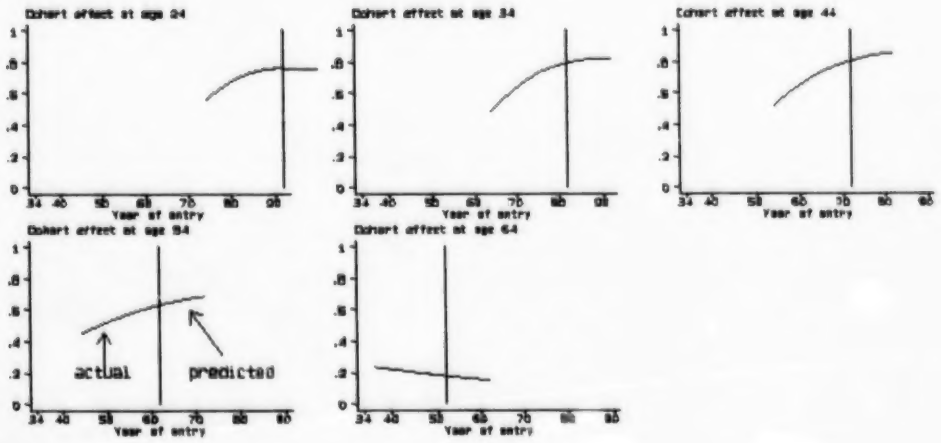

Figure 8: Predicted and actual cohort effects
Participation rate

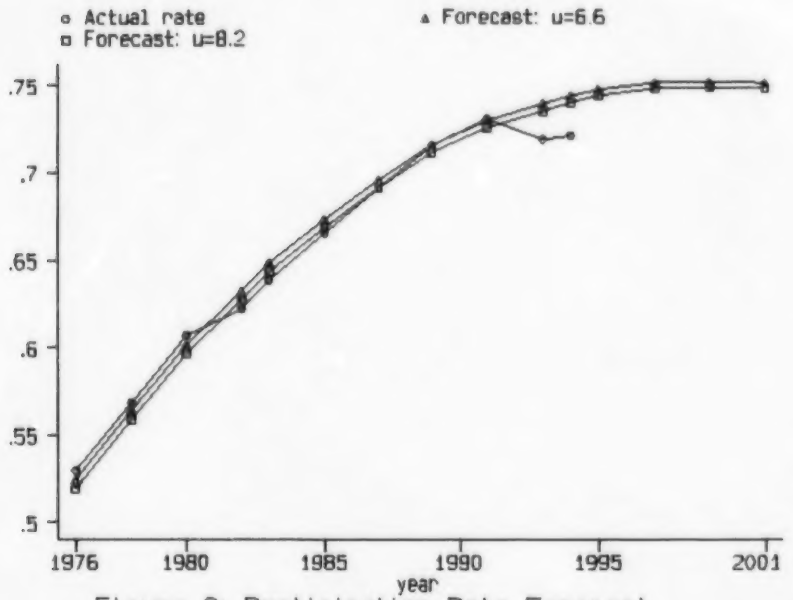

Figure 9: Participation Rate Forecast

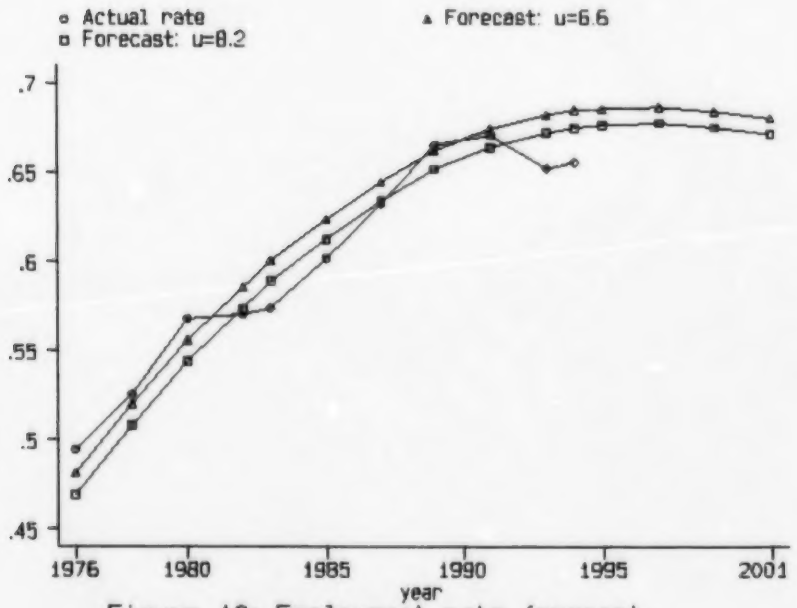

Figure 10: Employment rate forecast

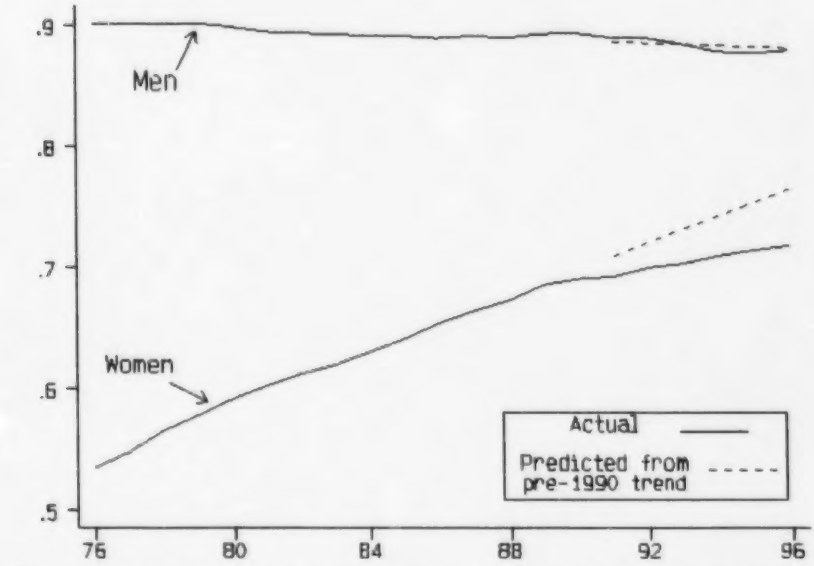

Figure 11: Participation Rates in the United States
Men and Women Age 25 to 64

References

Beaudry, Paul, and David Green. *"Cohort Patterns" in Canadian Earnings: Assessing the Role of Skill Premia in Inequality Trends*. National Bureau of Economic Research. Working Paper No. 6132, August 1997.

Beaudry, Paul, and Thomas Lemieux. *L'évolution du taux d'activité des femmes au Canada, 1976-1994 : Une analyse de cohortes*. Centre for Interuniversity Research and Analysis on Organizations. Working Paper No. 98s-06, March 1998.

Card, David, and W. Craig Riddell. "A Comparative Analysis of Unemployment in Canada and the United States," in D. Card and R. Freeman (Eds.) *Small Differences that Matter: Labor Markets and Income Maintenance in Canada and the United States*. Chicago: University of Chicago Press for National Bureau of Economic Research, 1993, pp. 149-189.